On the exponential convergence of the Kaczmarz algorithm

Liang Dai and Thomas B. Schön

Department of Information Technology, Uppsala University,

751 05 Uppsala, Sweden. E-mail: {liang.dai, thomas.schon}@it.uu.se

arXiv:1411.4017v2 [cs.SY] 10 Mar 2015

Abstract

The Kaczmarz algorithm (KA) is a popular method for solving a system of linear equations. In this note we derive a new exponential convergence result for the KA. The key allowing us to establish the new result is to rewrite the KA in such a way that its solution path can be interpreted as the output from a particular dynamical system. The asymptotic stability results of the corresponding dynamical system can then be leveraged to prove exponential convergence of the KA. The new bound is also compared to existing bounds.

Index Terms

Kaczmarz algorithm, Stability analysis, Cyclic algorithm.

I. Problem Statement

In this note, we discuss the exponential convergence property of the Kaczmarz algorithm (KA) [1]. Since its introduction, the KA has been applied in many different fields and many new developments are reported[2]-[13]. The KA is used to find the solution to the following system of *consistent* linear equations

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{1}$$

This work was supported by the project *Probabilistic modelling of dynamical systems* (Contract number: 621-2013-5524) funded by the Swedish Research Council.

where $\mathbf{x} \in \mathbb{R}^n$ denotes the unknown vector, $\mathbf{A} \in \mathbb{R}^{m\times n}, m \geq n$, rank$(\mathbf{A}) = n$ and $\mathbf{b} \in \mathbb{R}^m$. Define the hyperplane H_i as

$$H_i = \{\mathbf{x}|\mathbf{a}_i^T\mathbf{x} = b_i\},$$

where the i-th row of $\mathbf{A}$ is denoted by $\mathbf{a}_i^T$ and the i-th element of $\mathbf{b}$ is denoted by b_i. Geometrically, the KA finds the solution by projecting (or approximately projecting) onto the hyperplanes cyclically from an initial approximation $\mathbf{x}_0$, which reads as

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \lambda\frac{b_{i(k)} - \mathbf{a}_{i(k)}^T\mathbf{x}_k}{\|\mathbf{a}_{i(k)}\|_2^2}\mathbf{a}_{i(k)}, \tag{2}$$

where $i(k) = \text{mod}(k, m) + 1$. In the update equation (2), λ is the relaxation parameter, which satisfies $0 < \lambda < 2$. We use the Matlab convention mod$(\cdot,\cdot)$ to denote the *modulus after division* operation and $\|\cdot\|_2$ to denote the spectral norm of a matrix.

It is well-known that the KA is sometimes rather slow to converge. This is especially true when several consecutive row vectors of the matrix $\mathbf{A}$ are in some sense "close" to each other. In order to overcome this drawback, the Randomized Karczmarz Algorithm (RKA) algorithm was introduced in [4] for $\lambda = 1$. The key of the RKA is that, instead of performing the hyperplane projections cyclically in a deterministic order, the projections are performed in a random order. More specifically, at time k, select a hyperplane H_p to project with probability $\frac{\|\mathbf{a}_p\|_2^2}{\|\mathbf{A}\|_F^2}$, for $p = 1, \cdots, m$. Note that $\|\cdot\|_F$ is used to denote the Frobenius norm of a matrix. Intuitively speaking, the involved randomization is performing a kind of "preconditioning" to the original matrix equations [6], resulting in a faster exponential convergence rate, as established in [4].

The specific and predefined ordering of the projections in the KA makes it challenging to obtain a tight bound of the convergence rate of the method. In [16], the authors build up the convergence rate of the KA by exploiting the Meany inequality [17], which works for the case $\lambda = 1$ in (2). [18], [20] also established convergence rates for the KA, for $\lambda \in (0, 2)$. In Section III, we will compare these results in more detail.

In this note, we present a different way to characterize the convergence property of the KA described in (2). The key underlying our approach is that we interpret the solution path of the KA as the output of a particular dynamical system. By studying the stability property of this related dynamical system, we then obtain new exponential convergence results for the KA. Related to this, it is interesting to note that the so-called Integral Quadratic Constraints (IQCs) has recently

been used in studying the convergence rate of first-order algorithms applied to solve general convex optimization problems [19].

The note will be organized as follows. In the subsequent section we make use of the sub-sequence $\{\mathbf{x}_{jm} - \mathbf{x}\}_{j=0}^{\infty}$ to enable the derivation of the new exponential convergence result. In Section III we discuss its connections and differences to existing results. Conclusions and ideas for future work are provided in Section IV.

II. The new convergence result

First, let us introduce the matrix $\mathbf{B} \in \mathbb{R}^{m\times n}$, for which the i-th row $\mathbf{b}_i^T$ is defined as $\mathbf{b}_i \triangleq \frac{\mathbf{a}_i}{\|\mathbf{a}_i\|_2}, i = 1, \cdots, m$. Furthermore, let $\mathbf{P}_i \triangleq \mathbf{b}_i\mathbf{b}_i^T$ for $i = 1, 2, \cdots, m$ and let $\theta_k \triangleq \mathbf{x}_k - \mathbf{x}$ for $k \geq 0$. Using this new notation allows us to rewrite (2) according to

$$\theta_{k+1} = (\mathbf{I} - \lambda\mathbf{P}_{i(k)})\theta_k, \tag{3}$$

which can be interpreted as a discrete time-varying linear dynamical system. Hence, this relation inspires us to study the KA by employing the techniques for analyzing the stability properties of time-varying linear systems, see e.g. [14], [15].

In what follows we will focus on analyzing the convergence rate of the sub-sequence $\{\|\theta_{jm}\|^2\}_{j=0}^{\infty}$. Given the fact that $i(k) = \text{mod}(k, m) + 1$, we have

$$\theta_{(j+1)m} = \left(\prod_{i=1}^{m}(\mathbf{I} - \lambda\mathbf{P}_i)\right)\theta_{jm} \triangleq \mathbf{M}_m\theta_{jm}.$$

The following theorem provides an upper bound on the spectral norm of $\mathbf{M}_m$.

Theorem 1: Let $\rho \triangleq \|\mathbf{M}_m\|_2$ and $0 < \lambda \leq 2$, then it holds that

$$\rho^2 \leq \rho_1 \triangleq 1 - \frac{\lambda(2-\lambda)}{(2+\lambda^2 m^2)\|\mathbf{B}^\dagger\|_2^2}, \tag{4}$$

where $\mathbf{B}^\dagger$ denotes the pseudo-inverse of the matrix $\mathbf{B}$.

Proof: Let $\mathbf{v}_0 \in \mathbb{R}^n$ be a vector satisfying $\mathbf{M}_m\mathbf{v}_0 = \rho\mathbf{v}_0$, $\|\mathbf{v}_0\|_2 = 1$ and let $\mathbf{v}_i = (\mathbf{I} - \lambda\mathbf{P}_i)\mathbf{v}_{i-1}$ for $i = 1, \cdots, m$. It follows that $\mathbf{v}_m = \mathbf{M}_m\mathbf{v}_0$ and $\|\mathbf{v}_m\|^2 = \rho^2$.

Notice that $\mathbf{P}_i^2 = \mathbf{P}_i$, so we have

$$(\mathbf{I} - \lambda\mathbf{P}_i)^2 = \mathbf{I} - (2\lambda - \lambda^2)\mathbf{P}_i,$$

for $i = 1, \cdots, m$. Hence it holds that

$$\begin{aligned}\|\mathbf{v}_i\|^2 &= \mathbf{v}_{i-1}^T(\mathbf{I} - \lambda\mathbf{P}_i)^2\mathbf{v}_{i-1} \\ &= \mathbf{v}_{i-1}^T(\mathbf{I} - \lambda(2-\lambda)\mathbf{P}_i)\mathbf{v}_{i-1} \\ &= \|\mathbf{v}_{i-1}\|^2 - \lambda(2-\lambda)\|\mathbf{P}_i\mathbf{v}_{i-1}\|^2,\end{aligned}$$

which in turn implies that

$$\lambda(2-\lambda)\sum_{i=1}^{m}\|\mathbf{P}_i\mathbf{v}_{i-1}\|^2 = \|\mathbf{v}_0\|^2 - \|\mathbf{v}_m\|^2 = 1 - \rho^2. \tag{5}$$

Also, for any $i \in \{1, \cdots, m\}$, we have that

$$\begin{aligned}&\|\mathbf{v}_i - \mathbf{v}_0\| \\ &= \left\|\sum_{k=1}^{i}(\mathbf{v}_k - \mathbf{v}_{k-1})\right\| = \lambda\left\|\sum_{k=1}^{i}\mathbf{P}_k\mathbf{v}_{k-1}\right\| \\ &\le \lambda\sum_{k=1}^{i}\|\mathbf{P}_k\mathbf{v}_{k-1}\| \le \lambda\sqrt{i}\sqrt{\sum_{k=1}^{i}\|\mathbf{P}_k\mathbf{v}_{k-1}\|^2} \\ &\le \sqrt{\lambda i}\sqrt{\lambda\sum_{k=1}^{m}\|\mathbf{P}_k\mathbf{v}_{k-1}\|^2}\end{aligned}$$

Together with (5), we get

$$\|\mathbf{v}_i - \mathbf{v}_0\|^2 \le \frac{\lambda i}{2-\lambda}(1-\rho^2). \tag{6}$$

Meanwhile, we have that

$$\begin{aligned}&\lambda\mathbf{v}_0^T\mathbf{B}^T\mathbf{B}\mathbf{v}_0 \\ &= \lambda\sum_{k=1}^{m}\mathbf{v}_0^T\mathbf{P}_k\mathbf{v}_0 = \lambda\sum_{k=1}^{m}\|\mathbf{P}_k\mathbf{v}_0\|^2 \\ &= \lambda\sum_{k=1}^{m}\|\mathbf{P}_k[\mathbf{v}_{k-1} + (\mathbf{v}_0 - \mathbf{v}_{k-1})]\|^2 \\ &\le 2\lambda\sum_{k=1}^{m}\|\mathbf{P}_k\mathbf{v}_{k-1}\|^2 + 2\lambda\sum_{k=1}^{m}\|\mathbf{P}_k(\mathbf{v}_{k-1} - \mathbf{v}_0)\|^2 \\ &\le 2\lambda\sum_{k=1}^{m}\|\mathbf{P}_k\mathbf{v}_{k-1}\|^2 + 2\lambda\sum_{k=1}^{m}\|\mathbf{v}_{k-1} - \mathbf{v}_0\|^2.\end{aligned} \tag{7}$$

Together with (5) and (6), we have that

$$\lambda \mathbf{v}_0^T \mathbf{B}^T \mathbf{B} \mathbf{v}_0 \leq \frac{2(1-\rho^2)}{2-\lambda} + 2\lambda \sum_{k=1}^{m} \frac{\lambda(k-1)}{2-\lambda}(1-\rho^2)$$

or equivalently

$$\lambda \mathbf{v}_0^T \mathbf{B}^T \mathbf{B} \mathbf{v}_0 \leq \frac{1-\rho^2}{2-\lambda} \left(2 + \lambda^2 m(m-1)\right), \tag{8}$$

hence it follows that

$$\rho^2 \leq 1 - \frac{\lambda(2-\lambda)\mathbf{v}_0^T \mathbf{B}^T \mathbf{B} \mathbf{v}_0}{2 + \lambda^2 m(m-1)}.$$

Since $\mathbf{v}_0^T \mathbf{B}^T \mathbf{B} \mathbf{v}_0 \geq \frac{1}{\|\mathbf{B}^\dagger\|_2^2}$, we conclude that

$$\rho^2 \leq 1 - \frac{\lambda(2-\lambda)}{(2 + \lambda^2 m(m-1)) \|\mathbf{B}^\dagger\|_2^2}. \tag{9}$$

Finally, notice that $m(m-1) \leq m^2$ holds for any natural number m, which concludes the proof. ■

Remark 1: Notice that in the proof of Theorem 1, (7) is the main approximation step , and a better approximation here will lead to an improvement of the bound.

The following corollary characterizes the convergence of the KA under $\lambda = 1$, which will be used in the subsequent section to enable comparison to the results given in [16], [17]. We omit the proof since it is a direct implication of Theorem 1.

Corollary 1: For the KA with $\lambda = 1$ in (2), if $m \geq n \geq 2$, we have that

$$\rho^2 \leq 1 - \frac{1}{2m^2 \|\mathbf{B}^\dagger\|_2^2}. \tag{10}$$

Next, we will derive an improvement over the bound (4), enabled by partitioning the matrix $\mathbf{A}$ into non-overlapping sub-matrices. Let $q = \lceil \frac{m}{n} \rceil + 1$, where $\lceil x \rceil$ denotes the smallest number which is greater or equal to x. Define the following sets as $T_i = \{(i-1)n + 1, \cdots, in\}$, for $i = 1, \cdots, q-1$ and $T_q = \{(q-1)n + 1, \cdots, m\}$. Further, for $i = 1, \cdots, q$, define $\mathbf{B}_i$ as the sub-matrix of $\mathbf{B}$ with the rows indexed by the set T_i, and $\mathbf{N}_i = \prod_{j \in T_i} (\mathbf{I} - \lambda \mathbf{P}_j)$.

Corollary 2: Based on the previous definitions, and further assume that all the sub-matrices $\mathbf{B}_i$ for for $i = 1, \cdots, q$ are of rank n, then we have that

$$\rho^2 \leq \rho_2 \triangleq \prod_{i=1}^{q} \left(1 - \frac{\lambda(2-\lambda)}{(2 + \lambda^2 n(n-1)) \|\mathbf{B}_i^\dagger\|_2^2}\right) \tag{11}$$

Proof: Notice that since

$$\mathbf{M}_m = \mathbf{N}_q \mathbf{N}_{q-1} \cdots \mathbf{N}_2 \mathbf{N}_1,$$

we have that

$$\rho^2 = \|\mathbf{M}_m\|_2^2 \leq \prod_{i=1}^{q} \|\mathbf{N}_i\|_2^2. \tag{12}$$

For each $\mathbf{N}_i$, the spectral norm can be bounded analogously to what was done in Theorem 1, resulting in

$$\|\mathbf{N}_i\|_2^2 \leq 1 - \frac{\lambda(2-\lambda)}{(2+\lambda^2 n(n-1))\,\|\mathbf{B}_i^\dagger\|_2^2}.$$

Finally, inserting this inequality into (12) concludes the proof. ■

III. Discussion and numerical illustration

In Section III-A and III-B we compare our new bound with the bounds provided by the Meany inequality [16], [17] and the RKA, respectively. In Section III-B we also provide a numerical illustration. Section III-C is devoted to a comparison with the bound provided in [18], and finally Section III-D compares with the result given by [20].

A. Comparison with the bound given by Meany inequality

In the following, we assume that $m = n$ and $\lambda = 1$. Denote the singular values of $\mathbf{B}$ as $\sigma_1 \geq \sigma_2 \cdots \geq \sigma_n$, then the bound in [16], [17] given by the Meany inequality can be written as $\rho^2 \leq 1 - \prod_{i=1}^{n} \sigma_i^2$, and the bound given in (10) can be written as $\rho^2 \leq 1 - \frac{\sigma_n^2}{2n^2}$. This implies that when

$$\frac{\sigma_n^2}{2n^2} \geq \prod_{i=1}^{n} \sigma_i^2 \;\text{ i.e. }\; \prod_{i=1}^{n-1} \sigma_i^2 \leq \frac{1}{2n^2}, \tag{13}$$

holds, the bound in (10) is tighter. In the following lemma, we derive a sufficient condition, under which the inequality (13) holds.

Lemma 1: If $\sigma_{n-1}^2 \leq \frac{(n-2)^{n-2}}{2n^n}$ holds, the inequality in (13) is satisfied.

Proof: Notice that

$$\begin{aligned}\prod_{i=1}^{n-1} \sigma_i^2 = \left(\prod_{i=1}^{n-2} \sigma_i^2\right) \sigma_{n-1}^2 &\leq \left(\frac{\sum_{i=1}^{n-2} \sigma_i^2}{n-2}\right)^{n-2} \sigma_{n-1}^2 \\ &\leq \left(\frac{n}{n-2}\right)^{n-2} \sigma_{n-1}^2. \end{aligned} \tag{14}$$

The inequality (14) holds since

$$\sum_{i=1}^{n-2}\sigma_i^2 \leq \sum_{i=1}^{n}\sigma_i^2 = \|\mathbf{B}\|_F^2 = n.$$

Hence, if

$$\left(\frac{n}{n-2}\right)^{n-2}\sigma_{n-1}^2 \leq \frac{1}{2n^2}$$

holds, or equivalently if

$$\sigma_{n-1}^2 \leq \frac{(n-2)^{n-2}}{2n^n}$$

holds, then (13) holds, which concludes the proof. ■

Remark 2: Notice that the right hand side of the inequality in Lemma 1 is in the order of $\frac{1}{n^2}$ for large n. Another difference is that the bound provided by Theorem 1 depends explicitly on the size of matrix, while the bound provided by the Meany inequality does not.

B. Comparison with the bound given by the RKA

Let us now compare our new results to the results available for the RKA. Note that in this case, we set $\lambda = 1$. If $\{\theta_{jm}\}_{j=0}^{\infty}$ denotes the sequence generated by the RKA, then it holds that [4]

$$\mathbb{E}\|\theta_{jm}\|^2 \leq \left(1 - \frac{1}{\|\mathbf{A}\|_F^2\|\mathbf{A}^\dagger\|_2^2}\right)^{jm}\|\theta_0\|^2, \tag{15}$$

for $j \geq 1$, where $\mathbb{E}$ denotes the expectation operator with respect to the random operations up to index jm.

To compare (15) and (10), we make the assumption that $\mathbf{A}$ is a matrix with each row normalized, i.e. $\mathbf{A} = \mathbf{B}$, for simplicity. It follows that $\|\mathbf{B}\|_F^2 = m \leq m^2$, and

$$1 - \frac{1}{2m^2\|\mathbf{B}^\dagger\|_2^2} \geq 1 - \frac{1}{\|\mathbf{B}\|_F^2\|\mathbf{B}^\dagger\|_2^2}. \tag{16}$$

Furthermore, since $\|\mathbf{B}\|_F^2\|\mathbf{B}^\dagger\|_2^2 \geq 1$, we have that

$$1 - \frac{1}{\|\mathbf{B}\|_F^2\|\mathbf{B}^\dagger\|_2^2} \geq \left(1 - \frac{1}{\|\mathbf{B}\|_F^2\|\mathbf{B}^\dagger\|_2^2}\right)^m, \tag{17}$$

and combining (16) and (17), results in

$$1 - \frac{1}{2m^2\|\mathbf{B}^\dagger\|_2^2} \geq \left(1 - \frac{1}{\|\mathbf{B}\|_F^2\|\mathbf{B}^\dagger\|_2^2}\right)^m.$$

The above inequality implies that the bound given by (10) is more conservative than the one given by the RKA.

Next, a numerical illustration is implemented to compare the bounds given by (10), (11) and (15). The setup is as follows. Let $m = 30$ and $n = 3$, generate $\mathbf{A} = \text{randn}(30, 3)$ and normalize each row to obtain $\mathbf{B}$, generate $\mathbf{x} = \text{randn}(3, 1)$ and compute $\mathbf{y} = \mathbf{B}\mathbf{x}$. In the implementation of the RKA, we run $1\,000$ realizations with the same initial value $\mathbf{x}_0$ to obtain an average performance result, which is reported in Fig. 1.

From the left panel in Fig. 1, we can see that the bound (15) for characterizing the convergence of the RKA is closer to the real performance of the RKA, while the bounds given by (10) and (11) for bounding the convergence of the KA are further away from the real performance of the KA.

The right panel in Fig. 1 shows a zoomed illustration of the bound given by (10) and (11). We can observe that the bound given by (11) improves upon (10), which is enabled by the partitioning of the rows of the matrix.

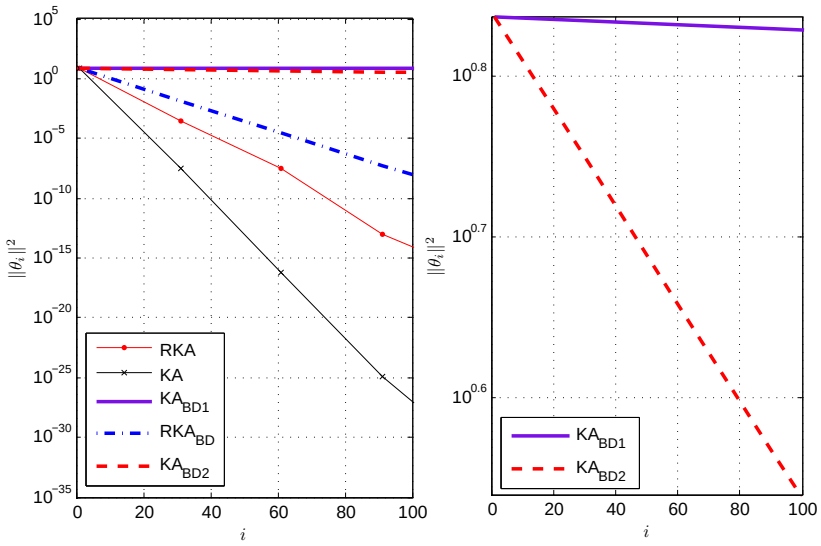

Fig. 1. In the left panel, the curves with tags 'KA' and 'RKA' illustrate the real performance of the KA and the RKA. The curves with tags 'KA_{BD1}', 'KA_{BD2}' and 'RKA_{BD}' illustrate the bounds given by (10), (11) and (15), respectively. In the right panel, a zoomed illustration for the curves 'KA_{BD1}' and 'KA_{BD2}' in the left panel is given.

C. Comparison with the bound given in [18]

To compare the result given in [18], we assume that $\mathbf{A} = \mathbf{B}$, and that they are square and invertible matrices. Under these assumptions, the involved quantity μ in Corollary 4.2 of [18] can be approximated by $\frac{1}{\sqrt{m}\|\mathbf{B}^\dagger\|_2}$ given the results in Theorem 2.2 of [18]. Hence, the convergence rate of the KA given by Theorem 3.1 in [18] can be written as

$$\rho^2 \leq 1 - \frac{\lambda(2-\lambda)}{m\left[1+(m-1)\lambda^2\right]\|\mathbf{B}^\dagger\|_2^2}, \tag{18}$$

where $\lambda \in (0,2)$. The result of the current work reads as

$$\rho^2 \leq 1 - \frac{\lambda(2-\lambda)}{(2+\lambda^2 m^2)\|\mathbf{B}^\dagger\|_2^2}, \tag{19}$$

where $\lambda \in (0,2)$. A closer look at the two bounds (18) and (19) reveals the following:

1) The optimal choice for the right hand side (RHS) of (18) is $\lambda = \frac{\sqrt{4m-3}-1}{2(m-1)}$, resulting in $\rho^2 \leq 1 - \frac{2}{m(\sqrt{4m-3}+1)\|\mathbf{B}^\dagger\|_2^2}$. When m is large, ρ^2 decreases with the speed $\frac{1}{m^{1.5}\|\mathbf{B}^\dagger\|_2^2}$.
2) When $\lambda = \frac{\sqrt{2}}{m}$ (a suboptimal choice for simplicity), (19) gives that $\rho^2 \leq 1 - \frac{\sqrt{2}(2-\frac{\sqrt{2}}{m})}{4m\|\mathbf{B}^\dagger\|_2^2}$. When m is large, ρ^2 decreases in the speed of $\frac{1}{m\|\mathbf{B}^\dagger\|_2^2}$, faster than the one in [18]. A comparison of both bounds when the optimal λ are chosen is given in Fig. (2).
3) When λ is chosen to be 1, both bounds decrease in the order of m^{-2}.

D. Comparison with the bound given in [20]

We will once again assume that each row of $\mathbf{A}$ is normalized, i.e. $\mathbf{B} = \mathbf{A}$. In [20] the authors makes use of a *subspace correction method* in studying the convergence speed of the KA. They show that (see eq. (31) in [20]), when the best relaxation parameter λ is chosen, ρ^2 can be bounded from above according to

$$\rho^2 \leq 1 - \frac{1}{\lfloor \log_2(2m) \rfloor \|\mathbf{B}\|_2^2 \|\mathbf{B}^\dagger\|_2^2}. \tag{20}$$

As we discussed in the previous section, when a near-optimal λ (i.e. λ is chosen as $\frac{\sqrt{2}}{m}$) is used, the upper bound implied by our analysis gives that

$$\rho^2 \leq 1 - \frac{\sqrt{2}(2-\frac{\sqrt{2}}{m})}{4m\|\mathbf{B}^\dagger\|_2^2}. \tag{21}$$

By assumption we have $\|\mathbf{B}\|_F^2 = m$, which implies that

$$\|\mathbf{B}\|_2^2 = \|\mathbf{B}^T\mathbf{B}\|_2 \geq \frac{\text{tr}(\mathbf{B}^T\mathbf{B})}{n} = \frac{\|\mathbf{B}\|_F^2}{n} = \frac{m}{n}. \tag{22}$$

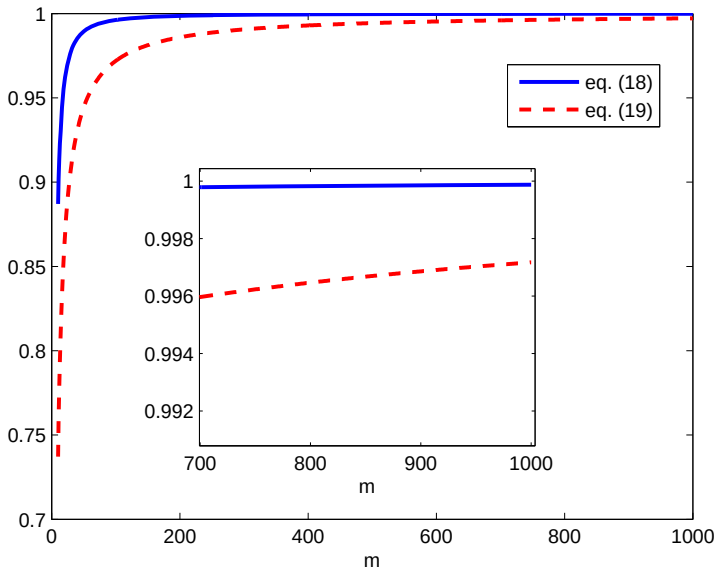

Fig. 2. The bounds (18) and (19) (using the optimal parameters) are plotted for the given $\|\mathbf{B}^{\dagger}\|_2 = 0.5$ and m ranging from 10 to 1 000. The result shows that the bound proposed in this work is always lower than the one given in [18] under the experimental settings.

Hence, the bound obtained by [20] will decrease with a speed of $\frac{1}{m\log_2(m)\|\mathbf{B}^{\dagger}\|_2^2}$ as m increases, while the present work gives the decreasing speed of $\frac{1}{m\|\mathbf{B}^{\dagger}\|_2^2}$.

IV. Summary

By studying the stability property of a time-varying dynamical system that is related to the KA we have been able to establish some new results concerning the convergence speed of the algorithm. The new results are also compared to several related, previously available results. Let us end the discussion by noting that the following two ideas can possibly lead to further improvements of the results. One potential idea is trying to improve the inequality in (7), since this part introduces much of the approximations in establishing the main result of the note; another idea is to try to find an optimal partitioning of the rows of the matrix $\mathbf{A}$, such that the right hand side of (11) is minimized.

V. ACKNOWLEDGEMENT

The authors would like to thank the reviewers for useful comments and pointing out the references [18], [20], and Marcus Björk for helpful suggestions.

REFERENCES

[1] S. Kaczmarz, Angenaherte Auflosung von Systemen linearer Gleichungen, Bulletin International de l'Académie Polonaise des Sciences et des Lettres, 35, 355-357, 1937.

[2] P. P. B. Eggermont, G. T. Herman, A. Lent, Iterative algorithms for large partitioned linear systems, with applications to image reconstruction, Linear algebra and its applications, 40, 37-67, 1981.

[3] Y. Censor, Row-action methods for huge and sparse systems and their applications, SIAM review, 23(4), 444-466, 1981.

[4] T. Strohmer, R. Vershynin, A randomized Kaczmarz algorithm with exponential convergence, Journal of Fourier Analysis and Applications, 15(2), 262-278, 2009.

[5] F. Natterer, The Mathematics of Computerized Tomography, Wiley, New York, 1986.

[6] T. Strohmer, R. Vershynin, Comments on the randomized Kaczmarz method, Journal of Fourier Analysis and Applications, 15(4), 437-440, 2009.

[7] D. Needell, Randomized Kaczmarz solver for noisy linear systems, BIT Numerical Mathematics, 50(2), 395-403, 2010.

[8] Y. Eldar, D. Needell, Acceleration of randomized Kaczmarz method via the Johnson-Lindenstrauss lemma, Numerical Algorithms, 58(2), 163-177, 2011.

[9] X. Chen, A. Powell, Almost sure convergence for the Kaczmarz algorithm with random measurements, Journal of Fourier Analysis and Applications, 18(6), 1195-1214, 2012.

[10] A. Zouzias, N. M. Freris, Randomized extended Kaczmarz for solving least-squares, SIAM Journal on Matrix Analysis and Applications, 2012.

[11] G. Thoppe, V. Borkar, D. Manjunath, A stochastic Kaczmarz algorithm for network tomography, Automatica, 2013.

[12] D. Needell, J. A. Tropp, Paved with Good Intentions: Analysis of a Randomized Block Kaczmarz Method, Linear Algebra and its Applications, 441, 199-221, 2014.

[13] L. Dai, M. Soltanalian, and K. Pelckmans, On the Randomized Kaczmarz Algorithm, IEEE Signal Processing Letters, 21(3), 330-333, 2014.

[14] L. Guo, Time-varying stochastic systems: stability, estimation and control, Jilin Science and Technology Press, Changchun, 1993.

[15] L. Guo, L. Ljung, Exponential stability of general tracking algorithms, IEEE Transactions on Automatic Control, 40(8), 1376-1387, 1995.

[16] A. Galántai, On the rate of convergence of the alternating projection method in finite dimensional spaces, Journal of mathematical analysis and applications, 310(1), 30-44, 2005.

[17] R. K. Meany, A matrix inequality, SIAM journal of Numerical Analysis, 6(1), 104-107, 1969.

[18] J. Mandel, Convergence of the cyclical relaxation method for linear inequalities, Mathematical programming, 30(2), 218-228, 1984.

[19] L. Lessard, B. Recht, A. Packard, Analysis and design of optimization algorithms via integral quadratic constraints, arXiv preprint, arXiv:1408.3595, 2014.

[20] P. Oswald, W. Zhou, Convergence Estimates for Kaczmarz-Type Methods, preprint, 2015.